Preprint typeset in JINST style - HYPER VERSION

The LHCb Trigger and its Performance

R. Aaij[o], J. Albrecht[m], F. Alessio[m], S. Amato[a], E. Aslanides[d], I. Belyaev[i], M. van Beuzekom[o], E. Bonaccorsi[m], R. Bonnefoy[c], L. Brarda[m], O. Callot[e], M. Cattaneo[m], H. Chanal[c], M. Chebbi[m], X. Cid Vidal[l], M. Clemencic[m], J. Closier[m], V. Coco[o], J. Cogan[d], O. Deschamps[c], H. Dijkstra[m]*, C. Drancourt[b], R. Dzhelyadin[j], M. Frank[m], M. Gandelman[a], C. Gaspar[m], V.V. Gligorov[m], C. Göbel[v], L.A. Granado Cardoso[m], Yu. Guz[j], C. Haen[m], J. He[e], E. van Herwijnen[m], R. Jacobsson[m], B. Jost[m], T.M. Karbach[m], U. Kerzel[m], P. Koppenburg[o], G. Krocker[f], C. Langenbruch[m], I. Lax[g], R. Le Gac[d], R. Lefèvre[c], J. Lefrançois[e], O. Leroy[d], L. Li Gioi[c], G. Liu[m], F. Machefert[e], I.V. Machikhiliyan[b,i], M. Magne[c], G. Mancinelli[d], U. Marconi[g], A. Martín Sánchez[e], M.-N. Minard[b], S. Monteil[c], N. Neufeld[m], V. Niess[c], S. Oggero[o], A. Pérez-Calero Yzquierdo[k], P. Perret[c], M. Perrin-Terrin[d], B. Pietrzyk[b], A. Puig Navarro[n], G. Raven[p], P. Robbe[e], H. Ruiz[k], M.-H. Schune[e], R. Schwemmer[m], J. Serrano[d], I. Shapoval[q,m], T. Skwarnicki[u], B. Souza De Paula[a], P. Spradlin[r], S. Stahl[f], V.K. Subbiah[m], S. T'Jampens[b], F. Teubert[m], C. Thomas[t], M. Vesterinen[m], M. Williams[s,1], M. Witek[h], A. Zvyagin[m]

[a] *Universidade Federal do Rio de Janeiro (UFRJ), Rio de Janeiro, Brazil*
[b] *LAPP, Université de Savoie, CNRS/IN2P3, Annecy-Le-Vieux, France*
[c] *Clermont Université, Université Blaise Pascal, CNRS/IN2P3, LPC, Clermont-Ferrand, France*
[d] *CPPM, Aix-Marseille Université, CNRS/IN2P3, Marseille, France*
[e] *LAL, Université Paris-Sud, CNRS/IN2P3, Orsay, France*
[f] *Physikalisches Institut, Ruprecht-Karls-Universität Heidelberg, Heidelberg, Germany*
[g] *Sezione INFN di Bologna, Bologna, Italy*
[h] *Henryk Niewodniczanski Institute of Nuclear Physics Polish Academy of Sciences, Kraków, Poland*
[i] *Institute of Theoretical and Experimental Physics (ITEP), Moscow, Russia*
[j] *Institute for High Energy Physics (IHEP), Protvino, Russia*
[k] *Universitat de Barcelona, Barcelona, Spain*
[l] *Universidad de Santiago de Compostela, Santiago de Compostela, Spain*
[m] *European Organization for Nuclear Research (CERN), Geneva, Switzerland*
E-mail: Hans.Dijkstra@cern.ch
[n] *Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland*
[o] *Nikhef National Institute for Subatomic Physics, Amsterdam, The Netherlands*
[p] *Nikhef National Institute for Subatomic Physics and VU University Amsterdam, Amsterdam, The Netherlands*
[q] *NSC Kharkiv Institute of Physics and Technology (NSC KIPT), Kharkiv, Ukraine*
[r] *School of Physics and Astronomy, University of Glasgow, Glasgow, United Kingdom*
[s] *Imperial College London, London, United Kingdom*
[t] *Department of Physics, University of Oxford, Oxford, United Kingdom*
[u] *Syracuse University, Syracuse, NY, United States*
[v] *Pontifícia Universidade Católica do Rio de Janeiro (PUC-Rio), Rio de Janeiro, Brazil, associated to [a]*
[1] *Massachusetts Institute of Technology, Cambridge, MA, United States*

ABSTRACT: The LHCb trigger and its performance based on data taken at the LHC in 2011 is presented. LHCb is designed to perform flavour physics measurements, and its trigger distinguishes charm and beauty decays from the light quark background. It uses a combination of lepton identification of particles, transverse momentum of particles and selects particles originating from hadrons which decay after a finite flight distance. The trigger reduces the ~ 11 MHz of bunch-bunch crossings with at least one non-elastic pp-interaction to 3 kHz of events which are written to storage in two trigger levels. The first level is implemented in hardware, while the next level is a software application which runs on all processors of a large computer farm. A data driven method is used to evaluate the performance of the trigger for several charm and beauty decay modes.

KEYWORDS: Trigger algorithms; Trigger concepts and systems (hardware and software).

*Corresponding author.

Contents

1. Introduction

LHCb is a single arm spectrometer with a pseudo-rapidity acceptance of $2 < \eta < 5$ optimised to perform flavour physics measurements at the LHC. The detector [1] layout is shown in figure 1. It consists of a silicon vertex detector surrounding the pp-interaction region (VELO), a silicon strip detector (TT), the dipole magnet, two Ring Imaging Cherenkov detectors (RICH 1&2), tracking detectors T1-T3, which consist of silicon strip detectors (IT) near the beam and straw tubes (OT) further out, a calorimeter system consisting of a Scintillating Pad detector (SPD), an electromagnetic calorimeter with pre-shower (ECAL, PS) and a hadronic calorimeter (HCAL), and muon chambers (M1-M5).

The LHCb trigger uses all of the above sub-systems, and its architecture has two levels, Level-0 (L0) and the High Level Trigger (HLT). L0 is a trigger implemented in hardware based on calorimeter and muon systems. L0 reduces the rate of crossings with interactions to below 1.1 MHz at which the whole detector can be read out. This maximum rate is imposed by the front-end (FE) electronics. The implementation of L0 is as described in reference [1], and hence its implementation will

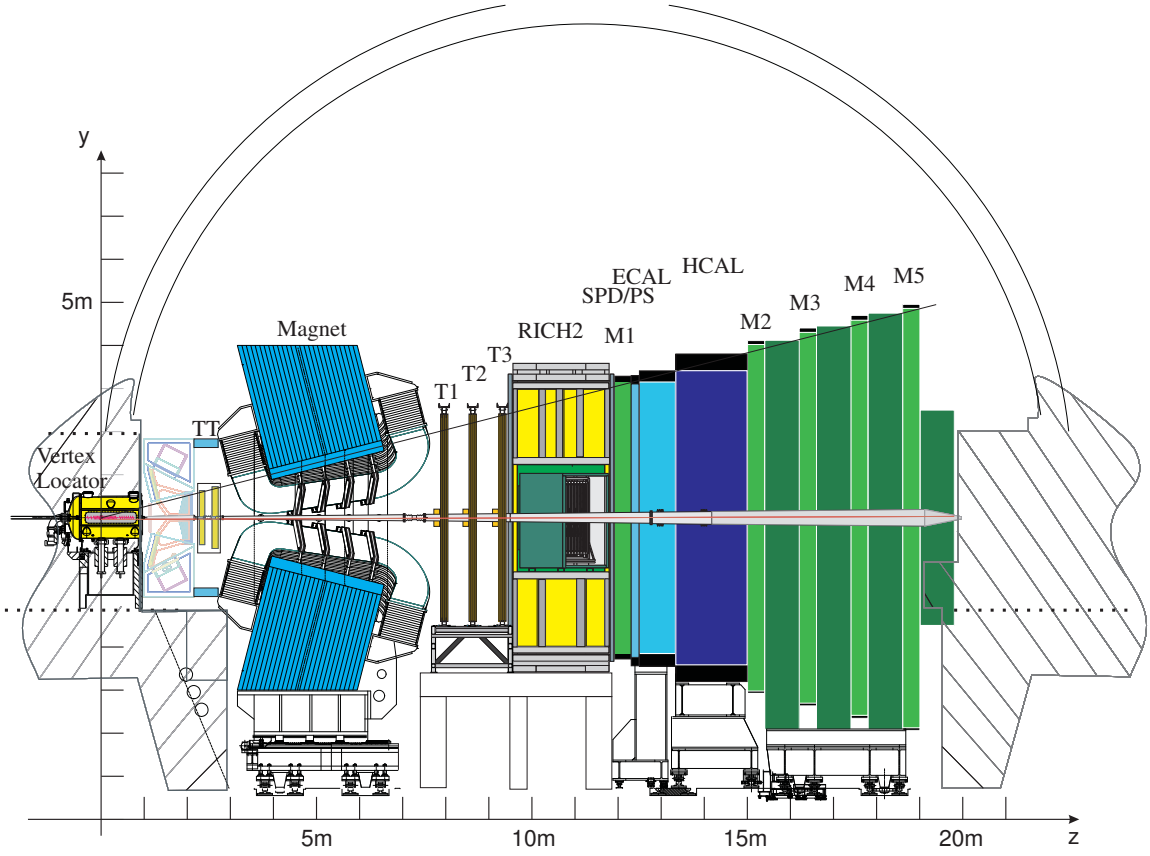

Figure 1. Layout of the LHCb detector.

only be described briefly in section 2. The HLT consists of a software application which runs on every processor of the Event Filter Farm (EFF). The actual HLT has evolved significantly compared to reference [1], which assumed that the LHC machine would operate with 25 ns bunch separation, and that LHCb would limit the number of visible pp-interactions[1] per bunch crossing to $\mu \approx 0.4$. However, the smallest bunch separation of the machine was 50 ns in the 2010 and 2011 physics runs, and in addition the LHCb detector performance was found not to be degraded up to $\mu \approx 2.5$. Hence, the HLT needed to adapt to running conditions rather different from those described in reference [1]. The HLT is subdivided in two stages, HLT1 and HLT2, which are described in detail in section 3. The HLT reduces the rate to ~ 3 kHz after which the events are written to storage. All events events written to storage are processed with more accurate alignment and calibration of the sub-detectors, and with reconstruction software which is more elaborate and allows for more redundancy than what can be used in the HLT. This part of the reconstruction and subsequent selection of interesting events will henceforth be referred to as off-line reconstruction and selection.

The method which is used to obtain a data driven determination of their performance is described in section 4. Section 5 will give the performance of the trigger in 2011 relative to off-line reconstruction and selection. Section 6 concludes with a summary of the trigger performance.

[1] A visible interaction is defined as having reconstructed at least two tracks in the VELO which point to the interaction envelope, and μ is the average number of visible interaction per bunch crossing.

2. Level-0 Trigger

L0 is divided into three independent triggers, the L0-Calorimeter trigger, L0-Muon trigger and the L0-PileUp trigger. The latter is not used to trigger on flavour physics events, but to aid the determination of the luminosity [4], and will not be further described in this paper. The L0 trigger system is fully synchronous with the 40 MHz bunch crossing signal of the LHC. The latencies are fixed and depend neither on the occupancy nor on the bunch crossing history. All L0 electronics is implemented in fully custom-designed hardware which make use of parallelism and pipe-lining to do the necessary calculations within the maximum latency of 4 μs. The trigger decisions are combined in a single L0-decision, which is transferred to the Readout Supervisor board (RS). The RS generates in addition a small rate of random (NoBias) triggers taking into account the bunch filling scheme of the machine[2]. The RS emulates the state of the FE-buffers to protect against their overflow. It also has information on the state of the buffers in the read-out boards of all sub-detectors and the availability of free nodes in the EFF. Based on this information it can accept or throttle a L0 trigger.

2.1 L0-Calorimeter Trigger Implementation

The L0-Calorimeter system uses information from SPD, PS, ECAL and HCAL. These four detectors are stacked along the beam axis (z-axis) and their longitudinal segmentation offers the possibility to distinguish between photon, electron and hadron showers. Transverse to the beam axis (x, y plane) the detectors are segmented into square cells. SPD, PS and ECAL are divided into three zones with ECAL cell rib sizes of 40.4 mm in the inner zone close to the beam pipe, 60.6 mm and 121.2 mm. HCAL is divided into two zones with rib size 131.3 mm and 262.6 mm. The SPD, PS and ECAL have the same geometry and are projective, i.e. the size of the cells in SPD and PS is adjusted to take into account the different z positions of the detectors. HCAL cells are larger but their boundaries always correspond to ECAL cell boundaries. The L0-Calorimeter system computes the transverse energy deposited in clusters of 2×2 cells using only cells of the same size. The transverse energy is defined as:

$$E_T = \sum_{i=1}^{4} E_i \sin\theta_i, \tag{2.1}$$

where E_i is the energy deposited in cell i and θ_i is the angle between the z-axis and a neutral particle assumed to be coming from the mean position of the interaction envelope hitting the centre of the cell. The ECAL and HCAL signals are read out and processed in FE boards (FEB) which cover an area of $(8+1) \times (4+1)$ cells, and the (+1) cells are shared between neighbouring FEB. Each FEB selects the highest E_T cluster among its 32 candidates. From these candidates, three types of candidates are built combining information as follows:

1. hadron candidate (`L0Hadron`): the highest E_T HCAL cluster. If there is a highest E_T ECAL cluster located in front of the cluster, the E_T of the hadron candidate is the sum of the E_T of the HCAL and ECAL clusters.

[2]Not all of the 3564 slots available for proton bunches around the machine are filled with protons. Most of the luminosity in 2011 was collected with 1296 bunches colliding in LHCb.

2. photon candidate (`L0Photon`): the highest E_T ECAL cluster with 1 or 2 PS cells hit in front of the ECAL cluster and no hit in the SPD cells corresponding to the PS cells. In the inner zone of the ECAL, an ECAL cluster with 3 or 4 PS cells hit in front of it is also accepted as photon. The E_T of the candidate is the E_T deposited in the ECAL alone.

3. electron candidate (`L0Electron`): same requirements as for a photon candidate, with in addition at least one SPD cell hit in front of the PS cells.

The E_T of the candidates is compared to a fixed threshold and events containing at least one candidate above threshold trigger L0. The total number of hits in the SPD is also determined, which is used to veto events which would take a disproportional large fraction of the available processing time in the HLT.

2.2 L0-Muon Trigger Implementation

The muon system consists of five muon stations (M1-M5), with pads in the high occupancy regions and horizontal and vertical strips elsewhere. Strips are combined to form logical pads for the muon trigger. The pad sizes are chosen to obtain projectivity towards the interaction region in the y-projection. Each quadrant of the muon detector is connected to a L0 muon processor, and there is no exchange of information between quadrants, hence muons traversing quadrant boundaries cannot be reconstructed in the trigger. Each of the four L0 muon processors look for the two muon tracks with the largest and second largest p_T in their quadrant. It searches for hits defining a straight line through the five muon stations and pointing towards the interaction point in the y-projection. In the x-projection the search is limited to muons with a $p_T \gtrsim 0.5$ GeV/c transverse to the beam line. The position of a track in the first two stations allows the determination of its p_T with a measured momentum resolution of ~ 25 % relative to the momentum of off-line reconstructed muon tracks. The trigger sets a single threshold on either the largest ${p_T}^{\text{largest}}$ of the eight candidates (`L0Muon`), or a threshold on ${p_T}^{\text{largest}} \times {p_T}^{\text{2nd largest}}$ (`L0DiMuon`).

3. High Level Trigger

An event accepted by L0 is transported by the On-line system from the detector FE electronics to permanent storage. A detailed description of its implementation can be found in reference [1] and references therein. Each event is transported to one of the 15440 processors of the EFF, which is subdivided in 50 sub-farms.

The HLT is an application in C++, of which 26110 copies run on the EFF. The HLT is based on the same software [2] as used throughout LHCb. The off-line full event reconstruction and selection requires about 2 s per event. During 2011 the L0 trigger rate was about 870 kHz, which given the available resources in the EFF, limits the time per event in the HLT to $\sim$30 ms. The HLT uses, where time allows, the same reconstruction algorithms as used off-line, and in the following these algorithms will not be explained in detail. The emphasis will be on those algorithms, or modifications to off-line algorithms, which had to be made for allowing the HLT to select events within 30 ms.

The HLT consist of trigger lines, each line being optimised to cover a certain class of events of interest. Each trigger line is configured by a python script, which defines the basic reconstruction

steps for objects on which event selections can be based, for instance VELO tracks, or tracks identified as a muon. Each line contains the selection parameters and where applicable down scaling fractions. All lines operate independently. Their independence is checked by running all lines individually on a sample of events selected by the NoBias trigger, and making sure that the results are compatible with running all lines simultaneously. A combination of lines, together with a L0 configuration, form an unique trigger with its associated Trigger Configuration Key (TCK), which is encoded for every event in the raw data. The architecture of the HLT assures that while lines are independent, all algorithms are executed only once per event. Python scripts allow inspecting and comparing TCK configurations. During 2011 running, a typical TCK contained 23 L0 lines, 38 HLT1 lines and 131 HLT2 lines. The lines which cover the main physics goals of LHCb [3], and take most of the rate, and the reconstruction algorithms which they share are described below, while the corresponding selection parameters and their performance are given in section 5.

The remaining lines are for luminosity measurements, pre-scaled physics lines with looser cuts, lines which trigger on very low multiplicity events and lines which look for large transverse momentum jets. In addition the trigger contains lines to accept NoBias events, lines which monitor events with inconsistent raw data or other errors during the HLT processing, lines which allow the VELO to monitor the position of the pp-interaction envelope and lines selecting calibration and monitoring data for fast feedback on the quality of the data.

3.1 HLT1

The off-line VELO reconstruction software is fast enough to allow a full 3D pattern recognition for all events which enter the HLT. In the off-line VELO pattern recognition a second pass is made on unused hits to enhance the efficiency for tracks which point far away from the beam-line, but in the HLT this search is not executed. At the start of each LHC fill the mean position of the pp-interaction envelope is determined using VELO tracks. This mean position perpendicular to the beam ($\mathrm{PV}_{\mathrm{xy}}^{\mathrm{mean}}$) is measured to be stable to within a few μm per fill. The VELO tracks are used to construct vertices with at least 5 tracks originating from them, and those vertices within a radius of 300 μm of $\mathrm{PV}_{\mathrm{xy}}^{\mathrm{mean}}$ are considered to be primary vertices (PV).

While in the off-line pattern recognition all VELO tracks are considered to look for the corresponding hits in the tracking stations downstream of the magnet, HLT1 limits the execution time by selecting VELO tracks which have a larger probability to originate from signal decays. HLT1 lines which do not require muons select VELO tracks based on their smallest impact parameter (IP) to any PV. In addition cuts are applied to the quality of the VELO track based on the number of hits on a track and the difference between the number of hits assigned to the VELO track and the number of hits expected given the number of VELO sensors a track traverses.

For events triggered by `L0Muon` or `L0DiMuon` a fast muon identification is performed in HLT1 to select VELO tracks which are muon candidates as follows. For every VELO track a search window is defined in the M3 station by extrapolating the VELO track in a straight line. The magnet does not bend tracks in the vertical plane, and multiple scattering dominates the vertical size of the search window. A muon candidate is required to have a momentum of at least 6 GeV/c, hence the horizontal search window size corresponds to the deflection of a 6 GeV/c track. Hits found inside the search window are combined with the VELO track to form candidate tracks for

a search for additional muon hits in stations M2, M4 and M5. A candidate track is provisionally accepted if it contains at least one hit in addition to the M3 hit. As the final step of the algorithm, a linear χ^2 fit in the horizontal plane is performed and the χ^2/ndf is required to be less than 25. As soon as the first candidate is found, the algorithm stops and the VELO track is tagged as a muon candidate.

For selected VELO tracks their track-segment in the OT and IT-stations are reconstructed to determine their momentum (forward tracking). Imposing a minimum momentum and transverse momentum ($p,\ p_T$) in the forward tracking significantly reduces the search windows which have to be opened in the IT and OT tracking stations and consequently reduces the required processing time. The reconstructed tracks are fitted using a Kalman filter [5] based track fit to obtain its χ^2 and a full covariance matrix at the start of the track. This fit differs from the off-line reconstruction in the use of a simplified material geometry description, less iterations are performed in the fit and consequently a less sophisticated removal of outlier hits is performed. The invariant mass resolution of $J/\psi \to \mu^+\mu^-$ determined in the HLT is measured to be 3 % larger than the 14 MeV/c^2 mass resolution obtained off-line, which shows that the resolution of the track parameters obtained in the HLT is sufficiently close to the off-line to allow sharp cuts in IP, momentum and mass. For tracks tagged as muon candidates, the off-line muon identification algorithm [6] is applied to the tracks to purify the muon sample.

3.2 HLT2

HLT1 reduces the rate to a sufficiently low level to allow forward tracking of all VELO tracks in HLT2. In contrast to the off-line reconstruction, where an approach with two tracking algorithms is used, in HLT2 only the algorithm based on seeding the search with VELO tracks is employed, which leads to a lower tracking efficiency compared to off-line by $1-2$ % per track. In addition, only tracks with a $p > 5$ GeV/c and $p_T > 0.5$ GeV/c are reconstructed, again to limit the size of the search windows in the OT and IT tracking stations.

Muon identification in HLT2 is performed using the off-line muon identification algorithm, but now for all tracks from the forward tracking. Tracks are also associated to ECAL clusters to identify electrons.

A large share of the 3 kHz output rate of HLT2 is selected by so-called topological lines, which are designed to trigger on partially reconstructed *b*-hadron decays. These lines in principle cover all *b*-hadrons with at least two charged particles in the final state and a displaced decay-vertex. The topological lines take an inclusive approach to maximise efficiency, making them less dependent on reconstruction inefficiencies imposed by the minimum ($p,\ p_T$) requirements and loss due to the single, non redundant, reconstruction mentioned above. In the following two sections the topological lines are described in more detail.

While the topological lines target inclusive *b*-hadrons, dedicated exclusive lines have been implemented requiring all decay particles to have been reconstructed in HLT2. They mainly use narrow mass windows to reduce their rate. These exclusive lines either target prompt *c*-hadron production, or allow to trigger on hadronic *b*-hadron decays without the necessity to use lifetime biasing selections to reduce the rate, and are described in section 3.2.3.

3.2.1 Topological Lines

Tracks are selected based on their χ^2/ndf, IP and muon or electron identification, henceforth referred to as "Topo-Tracks". Not all of the *b*-hadron final state particles need to satisfy these criteria. The trigger is designed to allow for the omission of one or more final state particles when forming the trigger candidate. Each Topo-Track is assigned the K^+ mass, irrespective of its particle identification. N-body candidates are built as follows: two Topo-Tracks are combined to form a 2-body object. A maximum distance of closest approach (DOCA) is imposed for the object to either become a 2-body candidate or to become the seed for a 3-body object. A 3-body object is made by combining a 2-body object and another Topo-Track, another maximum DOCA is imposed for the object to either become a 3-body candidate or input to a 4-body candidate. This DOCA is calculated between the 2-body object and the additional Topo-Track, hence is not the maximum DOCA of the three Topo-Tracks. This sequence of DOCA selections enhances the efficiency of the topological lines on $B \to DX$ decays. A similar procedure is followed when making 4-body candidates from 3-body objects and an additional Topo-Track. If a n-body candidate only contains a subset of the daughter particles, its invariant mass (m) will be less than the mass of a *b*-hadron. Thus, a mass window would need to be very loose if the trigger is to be inclusive. Instead a corrected mass ($m_{\rm corr}$) is used which is defined as:

$$m_{\rm corr} = \sqrt{m^2 + |p'_{T\rm miss}|^2} + |p'_{T\rm miss}| \tag{3.1}$$

where $p'_{T\rm miss}$ is the missing momentum transverse to the direction of flight, which is defined by the PV and the n-body vertex [7]. In case of multiple PVs, the PV with respect to which the n-body combination has the smallest IP is used. The quantity $m_{\rm corr}$ would be the minimal mass of the parent if a mass-less particle was omitted from the trigger candidate. Prompt *c*-hadrons which are erroneously combined with another track constitute a significant fraction of the n-body candidates. These candidates are rejected by requiring that all (n-1)-body objects used by a n-body line either have a mass greater than 2.5 GeV/c^2 or that they have a significant IP to all PVs.

To select a n-body candidate, cuts are applied to the following variables: $\sum |p_T|$, p_T^{min}, m, $m_{\rm corr}$, DOCA, IPχ^2 and flight distance (FD) χ^2. It was found using NoBias events for background, and signal events from simulated (MC) events that a larger rejection power could be achieved for the same signal efficiency by combining the above variables in a multivariate selection.

3.2.2 Topological Multivariate Lines

To combine the variables mentioned above, a boosted decision tree (BDT) was chosen, which has already been successfully used elsewhere [8].

All multivariate classifiers select n-dimensional regions of a multivariate space by learning from the training samples provided to them. If selected regions are small relative to the resolution of the detector, the signal could oscillate between regions resulting in, at best, a less efficient trigger or, at worst, a trigger which is very difficult to understand. The topological trigger is designed to be inclusive; however, it is simply not possible to use every known *b*-hadron decay in the training. Hence, it is important to ensure that the BDT is learning common *b*-hadron decay traits and not a large sum of specific ones. Another concern is that a tree is an extremely large set of if/else statements and it can take a long time to evaluate the response for each event. A way to avoid

Table 1. MC signal samples used to train the BBDT, where K means $K^{\pm}$ and π means $\pi^{\pm}$.

parent	daughters
B^+	$K\pi\pi, D_{[K\pi]}\pi, D_{[hhhh]}K, D_{[K_S\pi\pi]}K, D_{[K\pi\pi]}K\pi$
B^0	$K^*_{[K\pi]}\mu\mu, K^*_{[K\pi]}ee, D_{[K\pi\pi]}\pi, K\pi, D_{[K\pi]}K\pi, D^*_{[D(K\pi)\pi]}\ mu\nu, D_{[K\pi\pi]}K\pi\pi$
B_s	$D_{s[KK\pi]}\pi, D_{s[KK\pi]}K\pi\pi, K^*_{[K\pi]}K^*_{[K\pi]}$
Λ_b	$\Lambda_{c[pK\pi]}\pi, \Lambda_{c[pK\pi]}K\pi\pi$

Table 2. Allowed split points in the BBDT. The variables are explained in the text.

variable	cuts(2, 3, 4-body)	allowed splits
$\sum \lvert p_T \rvert$ [GeV/c]	$> 3, 4, 4$	3.5, 4, 4.5, 5, 6, 7, 8, 9, 10, 15, 20
$p_T^{\min}$ [GeV/c]	> 0.5	0.6, 0.7, 0.8, 0.9, 1, 1.25, 1.5, 1.75, 2, 2.5, 3, 4, 5, 10
m [GeV/c^2]	< 7	2.5, 4.75
m_{corr} [GeV/c^2]		2, 3, 4, 5, 6, 7, 8, 9, 10, 15
DOCA [mm]	< 0.2	0.05, 0.1, 0.15
IPχ^2		20
FD$\chi^2/100$	> 1	2, 3, 4, 5, 6, 7, 8, 9, 10, 25, 50, 100

these issues is to discretise all of the variables, henceforth referred to as Bonsai BDT (BBDT). The BBDT enforces that the smallest interval that can be used satisfies $\Delta x_{\min} > \delta_x$ for all x values and on all leaves, where $\delta_x = \text{MIN}\{|x_i - x_j| : \ x_i, \ x_j \in x_{discrete}\}$. The constraints governing the choice of x_{discrete} are then as follows: (1) δ_x should be greater than the resolution on x and be large with respect to the expected variations in x and (2) the discretisation should reflect common b-hadron properties. The optimal discretisation scheme for each variable was determined by first training a BBDT with a very large number of discretisation values and then gradually decreasing this number as low as possible without losing much in performance. The training signal samples were MC events which contain as signal B^+, B^0, B_s or Λ_b decays[3] with decay modes as given in table 3.2.2, while the background sample was NoBias data recorded in 2010. Only MC events for which all products of the signal decay are reconstructed using the off-line reconstruction are used for the training. Table 3.2.2 shows the discretisation scheme for each of the variables used in the BBDT. The discretisation also allows converting the large number of if/else statements into a one-dimensional array of response values, and thus consuming only a small fraction of the available processing time.

3.2.3 Exclusive Lines

In the topological trigger described in the previous section there is an explicit veto on prompt charm. The selection of prompt charm decays is achieved by HLT2 lines which require an exclusive reconstruction of the decay products, and have tight cuts on the invariant mass of the reconstructed candidates. While HLT2 reconstruction and selection allows to obtain good efficiencies for b-hadrons, the $p_T > 0.5$ GeV/c constraint reduces the efficiency for the exclusive selection of

[3]Charge conjugate hadrons are always implied.

charm decays with more than two final state particles. To enhance the reconstruction efficiency for these exclusive lines, the lines first look for a two-prong secondary vertex. Selection cuts are imposed on the maximum invariant mass of the two tracks, the quality of their vertex, the sum of the transverse momenta of the tracks and $m_{\rm corr}$. These initial cuts reduce the rate sufficiently to allow for the forward tracking of the remaining VELO tracks, but now with a relaxed $p_T > 0.25$ GeV/c constraint and only using the hits in the tracking system which have not been used by the previous pass in forward tracking. Subsequently two-prong candidates are combined with other tracks, which now include the low p_T tracks, to form exclusively reconstructed candidates. The combinatorial background is reduced by tight requirements on the mass and on the angle between the momentum of the B and the vector connecting the PV with the B vertex. HLT2 contains 28 trigger lines dedicated to selecting prompt charm.

Another example of an exclusive trigger is a dedicated line selecting the decay $B_s \to K^+K^-$ while avoiding cuts which bias the B_s lifetime. This implies that a cut on IP cannot be used, while this is a powerful variables to reject combinatorial background. In order to enrich the B_s candidates two dedicated neural networks based on the NeuroBayes neural network package [9] are used. In a first step, kinematic constraints such as the transverse momentum of the final state particles and the helicity angle in the rest frame of the B_s candidate are used to reduce the rate. This allows to run the comparatively slow particle identification algorithm using the RICH sub-detector on the events selected by the first neural network. This information is then included in a second neural network which uses both kinematic and particle ID information to make the final selection.

4. Data Driven Trigger Performance Determination

The trigger performance is evaluated relative to the off-line reconstruction and selections, and thus contains only the additional inefficiency due to simplifications used in the trigger, possible alignment inaccuracies, worse resolution or harder cuts imposed by rate and/or processing time limitations. The off-line selections chosen here are representative of selections used in most analyses, and the channels used and selections applied are listed in Tables 4-4.

In what follows, the term "signal" refers to a combination of tracks which form the off-line reconstructed and selected *b* or *c*-hadron candidate. The trigger efficiency of events with a signal is determined from the data by classifying an event which would also have been triggered without the signal as a TIS (Trigger Independent of Signal) event, and/or classifying the event where the signal alone is sufficient to trigger the event as a TOS (Trigger On Signal) event. The trigger records all the information needed for such a classification. All strips, straws, cells and pads of the sub-detectors have an unique identifier, and these identifiers are written in a trigger report in the data stream for every line which accepts an event. Global event variables, such as primary vertices or the SPD multiplicity, are not considered in this classification. The criteria used to associate a trigger object with an off-line track are as follows:

- L0-Calorimeter: the off-line track is extrapolated to the *z*-position of the calorimeter (ECAL or HCAL), and the cell intersecting with the track and its eight neighbours are considered signal cells. If any of the 2×2 cells of a L0-Calorimeter cluster above the threshold coincides with a signal cell, this cluster is associated with the off-line track. If non of the cells overlap, the cluster is not associated with the off-line track.

Table 3. Off-line selection cuts applied for the modes $D^0 \to K^-\pi^+$ and $D^+ \to K^-\pi^+\pi^+$. The angle α is the angle between the momentum of the D and the vector connecting the PV with the D vertex.

Variable	$D^0 \to K^-\pi^+$	$D^+ \to K^-\pi^+\pi^+$
Track p_T [MeV/c]	> 800	> 250
Track Σp_T [MeV/c]	-	> 2800
Track χ^2 /ndf	< 3	< 5
Track p [MeV/c]	> 5000	> 2000
Track IPχ^2	> 9	> 4
$\geqq 2$ tracks IPχ^2	-	> 10
D p_T [MeV/c]	> 2000	> 1000
D IPχ^2	-	< 12
D FDχ^2	> 40	> 100
D cos(α)	> 0.9999	> 0

Table 4. Offline selection cuts applied for the modes $B^+ \to J/\psi K^+$ ($J/\psi \to \mu^+\mu^-$), $B^0 \to D^+\pi^-$ ($D^+ \to K^-\pi^+\pi^+$), and $B^+ \to D^0\pi^-$ ($D^0 \to K^-\pi^+$). The term "bachelor" refers to a track coming directly from the B decay. In addition, loose particle identification requirements are placed on all tracks, which do not significantly alter the momentum spectrum and are omitted for brevity. Mass windows are applied at approximately 3σ around the D and J/ψ masses to enhance purity. The angle α is the angle between the momentum of the B and the vector connecting the PV with the B vertex.

Variable	$J/\psi K^+$	$D^+\pi^-$	$D^0\pi^-$
Track χ^2 /ndf	< 4		
Track p [MeV/c]	> 2000		
Track p_T [MeV/c]	> 300	> 250	
Track IP χ^2	-	> 9	
Bachelor p_T [MeV/c]	> 1000	> 500	
Bachelor p [MeV/c]	-	> 5000	
Bachelor IP χ^2	-	> 16	
J/ψ (D) vertex χ^2	< 16	< 9	
J/ψ (D) DOCA χ^2	< 30	-	
J/ψ (D) p_T [MeV/c]	-	> 2000	
J/ψ (D) FD χ^2	-	> 100	
B p_T [MeV/c]	> 2000	> 1500	
B τ [ps]	> 0.2	> 0.2	
B IP χ^2	< 25	< 16	
B vertex χ^2	< 10	< 9	
B cos(α)	-	> 0.9999	
B FD χ^2	-	> 16	

- L0-Muon: the trigger records the M1, M2 and M3 hits used to form the L0 muon candidate. If at least two of the three hits are shared with an off-line reconstructed muon the L0 muon

Table 5. $B^0 \to J/\psi K^{*0}$ ($J/\psi \to e^+e^-$) offline selection cuts. The angle α is the angle between the momentum of the B and the vector connecting the PV with the B vertex.

Variable	$B^0 \to J/\psi K^{*0}$
Track χ^2	< 5
Sum of all track IP χ^2	> 50
$e^\pm$ IP χ^2	> 2.25
$p_T(e^\pm)$	$>$350 MeV/c
$p_T(K^\pm)$	> 400 MeV/c
$p_T(\pi^\pm)$	> 300 MeV/c
$p(K^\pm)$	> 3 MeV/c
$p(\pi^\pm)$	> 3 MeV/c
$K^\pm$ IP χ^2	> 4
$\pi^\pm$ IP χ^2	> 4
K^{*0} mass interval	766-1026 MeV/c^2
K^{*0} vertex χ^2	< 15
K^{*0} IP χ^2	> 1
J/ψ vertex χ^2	< 15
J/ψ FD	> 1 mm
J/ψ mass interval	2.2-4.2 GeV/c^2
B vertex χ^2	< 9
B cos(α)	> 0.999872

is associated with the off-line track. Non associated L0 muons have no hits overlapping between the L0 muon and the muon hits of the off-line track.

- HLT tracks: a HLT track has VELO hits and hits in the OT and/or the IT. In addition it can have TT hits and hits in the muon chambers M2-M5. A HLT track is not associated with an off-line track if it shares no hits with the off-line track. Associated tracks require that the fraction of HLT track hits which overlaps with the off-line track is at least 70 % in the VELO, 70 % in TT if applicable and 70 % of OT and IT combined. For muons the association requirement is that at least 60 % of the HLT muon hits overlap with the off-line muon.

Events classify as TOS if trigger objects which are associated with the signal are sufficient to trigger the event. Events classify as TIS if it could have been triggered by those trigger objects which are not associated to the signal. A number of events can be classified as TIS and TOS simultaneously ($N^{\text{TIS\&TOS}}$), which allows the extraction of the trigger efficiency relative to the off-line reconstructed events from data alone. The efficiency to trigger an event independently of the signal, ε^{TIS}, is given by $\varepsilon^{\text{TIS}} = N^{\text{TIS\&TOS}}/N^{\text{TOS}}$, where N^{TOS} is the number of events classified as TOS.

Table 6. $B^0 \to K^{*0}\gamma$ offline selection cuts. The angle α is the angle between the momentum of the B and the vector connecting the PV with the B vertex.

Variable	$B^0 \to K^{*0}\gamma$
Track χ^2	< 5
Track IP χ^2	> 25
Track p_T	$> 500\,\text{MeV}$
Max track p_T	$> 1200\,\text{MeV}$
$\gamma\, E_T$	> 2600 MeV
K^{*0} mass interval	846-946 MeV/c^2
K^{*0} vertex χ^2	< 9
B p_T	> 3000 MeV/c
B IP χ^2	< 9
B cos(α)	> 0.9998
B FD χ^2	> 100

The efficiency to trigger an event on the signal alone, ε^{TOS}, is given by $\varepsilon^{\text{TOS}} = N^{\text{TIS\&TOS}}/N^{\text{TIS}}$, where N^{TIS} is the number of events classified as TIS. The total trigger efficiency for events containing the signal can then be computed as $\varepsilon^{\text{TIS}} \times N^{\text{Trig}}/N^{\text{TIS}}$, where N^{Trig} is the total number of signal events.

The phase-space distribution of the signal is affected by the TIS requirement. This is illustrated in Fig 2, which shows the p_T distribution of $D^+ \to K^-\pi^+\pi^+$ selected from NoBias and TIS events. The p_T of TIS events is harder, which results in a larger signal efficiency integrated over all phase-

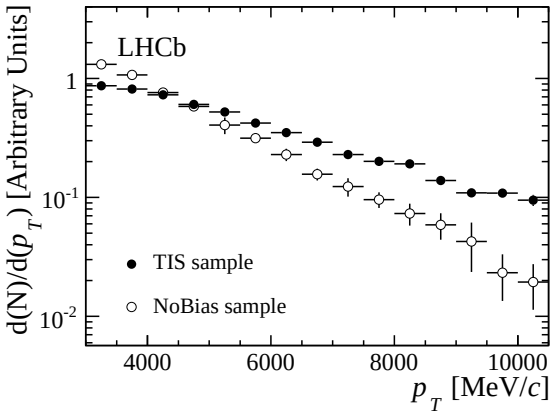

Figure 2. Comparison of the p_T distributions of $D^+ \to K^-\pi^+\pi^+$ selected in NoBias and TIS events.

space. This bias in phase-space can be understood considering that $c\bar{c}$, or $b\bar{b}$, pairs are correlated in phase-space at production, and TIS events are predominantly triggered by the decay products of the hadron which contains the other heavy quark. For charm decays TIS events could enhance the

non prompt charm component by triggering on the other decay products of a B to D decay. Hence, the trigger performance will be presented as a function of the signal p, p_T and lifetime (τ).

5. Trigger Performance

At each trigger level the different trigger lines need to compete for their share of the available resources. To determine the different selections for the trigger lines, a so-called bandwidth division procedure has been adopted, which is described in the next section. The performance of the different trigger lines with the thresholds as determined by the bandwidth division will be presented for L0, HLT1 and HLT2 for the channels listed in Tables 4-4.

In all off-line selected signal samples the level of background is significantly lower than the signal, however substantial differences in trigger efficiency are observed for true signal and background. Hence, all trigger performances are extracted by determining the signal component by fits to their invariant mass distributions.

5.1 Bandwidth Division Procedure

The bandwidth division minimises the overall loss in efficiency by minimising a χ^2 which is defined as:

$$\chi^2 = \sum_{\text{signal}} \sum_{\text{lines}} (1 - \frac{\varepsilon^{\text{signal, line}}}{\varepsilon^{\text{signal, line}}_{\text{max}}})^2 \quad , \tag{5.1}$$

where $\varepsilon^{\text{signal, line}}_{\text{max}}$ is the L0×HLT trigger efficiency for a signal with the full resources dedicated to that signal and a specific trigger line alone, and $\varepsilon^{\text{signal, line}}$ is obtained using one fixed set of selections for all signal channels and lines simultaneously. Cuts in all trigger lines are varied by running an emulation of L0 and executing the HLT application. This emulation includes FE-buffer overflow emulation and the available proocessing power in the EFF as throttle conditions.

For signal the following MC generated and off-line reconstructed and selected channels have been chosen to represent both the main physics goals of LHCb and to cover all the trigger components which need to be tuned: $B_s \to J/\psi(\mu^+\mu^-)\phi(K^+K^-)$, $B_s \to \mu^+\mu^-$, $B^0 \to K^{*0}\mu^+\mu^-$, $B_{(s,d)} \to \mu^+ X$, $D^{*+} \to D(\mu^+\mu^-)\pi^+$, $D^+ \to K^-\pi^+\pi^+$, $B^0 \to K^{*0}\gamma$, $B^0 \to K^+\pi^-$, $B^0 \to K^{*0}e^-e^+$, $B^+ \to K^+\pi^-\pi^+$, $B_s \to D_s^-(K^+K^-\pi^-)K^+$, $B^+ \to \bar{D}^0(K_S\pi^+\pi^-)K^+$. And from 2010 NoBias events at $\sqrt{s} = 7$ TeV selected $D^0 \to K^-\pi^+$, and 2010 NoBias events with $\mu = 1.4$ as background. All performance results are given for 1296 colliding bunches in LHCb, which corresponds to a bunch crossing rate with at least one visible pp-interaction of ~ 11 MHz. While the MC samples show a good agreement with data as far as the signal decay products is concerned, the multiplicity of hits in the trackers and SPD downstream of the magnet in MC is about ~ 30 % too low. To compensate for this the bandwidth division replaces the MC simulated SPD hit multiplicity with the SPD multiplicity extracted from real data (SPD_{mult}). In addition, the signal efficiencies are multiplied with a penalty function derived from MC and defined as $(1 - 0.026 \times \text{SPD}_{\text{mult}}/360)^n$, where n is the charged track multiplicity of the signal decay products. The bandwidth division yielded the following rates for NoBias events: 870 kHz for L0, 43 kHz for HLT1 and 3 kHz for HLT2.

5.2 L0 Performance

Events with a large occupancy in OT and IT consume a disproportionally large fraction of the available processing time in the HLT. The SPD multiplicity measured in L0 is a good measure of this multiplicity, which allows an early rejection of events which require a relatively large processing time. The bandwidth division has determined the SPD cut to be < 900 for events triggered by `L0DiMuon`, and < 600 for all other L0 triggers. Events with a SPD multiplicity larger than 600 consume four times more executon time in the HLT per event than the average consumption of events with a multiplicity below 600 SPD hits. The fraction of events rejected due to these cuts has been determined for prompt charm production with $D^+ \to K^-\pi^+\pi^+$ selected from NoBias events, and is found to be 7.4 ± 0.3 (0.05 ± 0.01) % for a cut on 600 (900) in SPD multiplicity. For b-hadron production this cannot be determined from NoBias events, since $\sigma_{b\bar{b}}$ times a typical branching ratio is too small to extract a large enough sample of signal events in NoBias events. The sample of events used to extract this SPD efficiency is $J/\psi \to \mu^+\mu^-$ events with $\tau_{J/\psi} > 0.3$ ps, selected by `L0DiMuon` to cover SPD multiplicities up to 900, and only one PV reconstructed. In this sample of 89830 signal events the mean SPD multiplicity is 227, with 2 events with a SPD multiplicity > 800. The SPD multiplicity of these events is convolved with the SPD multiplicity in NoBias events with $\mu = 1.4$. The fraction of events with a SPD multiplicity > 600 (900) is found to be 8.8 ± 0.6 (0.5 ± 0.2) %. All efficiencies quoted below are given relative to the sample after the SPD multiplicity cut.

Table 5.2 lists the cuts for the L0 lines. About 20 % of the events accepted by L0 are selected by more than one trigger line, giving a total L0 rate of 870 kHz prior to throttling. `L0Muon` is the main

Table 7. Trigger cuts of L0 lines and their rates prior to throttling. The definition of the lines is given in section 2.

L0 lines	Threshold	SPD multiplicity	Rate
`L0Muon`	$p_T > 1.48$ GeV/c	< 600	340 kHz
`L0DiMuon`	$\sqrt{p_T{}^{\text{largest}} \times p_T{}^{\text{2nd largest}}} > 1.296$ GeV/c	< 900	75 kHz
`L0Hadron`	$E_T > 3.5$ GeV	< 600	405 kHz
`L0Electron`	$E_T > 2.5$ GeV	< 600	160 kHz
`L0ElectronHi`	$E_T > 4.2$ GeV	< 600	27 kHz
`L0Photon`	$E_T > 2.5$ GeV	< 600	80 kHz
`L0PhotonHi`	$E_T > 4.2$ GeV	< 600	10 kHz

trigger for particle decays with one or more muons in the final state. `L0DiMuon` recovers part of the events with a SPD multiplicity > 600 for a small increase in rate. The performance of `L0Muon` and `L0DiMuon` are shown in figure 3 for $B^+ \to J/\psi K^+$ as a function of p_T (J/ψ). `L0DiMuon` increases the number of signal events by 4.9 %, of which 87 % have a SPD multiplicity larger than 600 hits. L0 requires a muon candidate to have a hit in all five muon stations, while off-line even down to two stations suffices to identify a muon. As a result `L0DiMuon` has a maximum efficiency of ~ 80 % even for a J/ψ with large p_T. `L0Muon` recovers this loss for lower SPD multiplicities

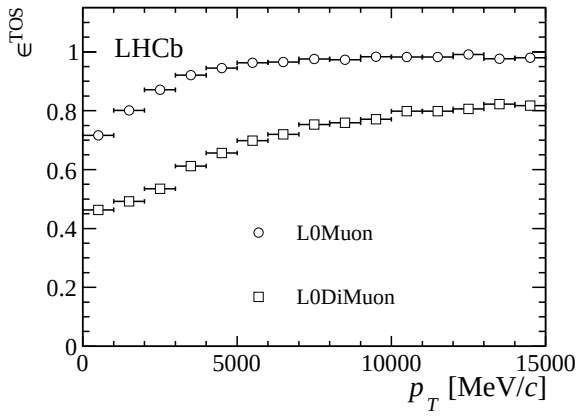

Figure 3. Efficiency $\varepsilon^{\rm TOS}$ of $B^+ \to J/\psi K^+$ ($J/\psi \to \mu^+\mu^-$) as a function of p_T (J/ψ) for `L0Muon` and `L0DiMuon` lines.

and decays with more muons at large p_T.

`L0Hadron` selects heavy flavour decays with hadrons in the final state. The performance of `L0Hadron` is shown in figure 4 for $B^0 \to D^+\pi^-$, $B^+ \to D^0\pi^-$, $D^0 \to K^-\pi^+$ and $D^+ \to K^-\pi^+\pi^+$ as a function of p_T and p of the signal. At low p_T `L0Hadron` has a better efficiency for b-hadrons than for c-hadrons due to the larger b-hadron mass. Once the p_T of the hadron is above the b-hadron mass, the decays with less multiplicity show the better efficiency. Since most of the hadrons are produced at small p_T, the efficiency as a function of p shows a better performance of `L0Hadron` for b-hadrons reflecting its better performance at low p_T.

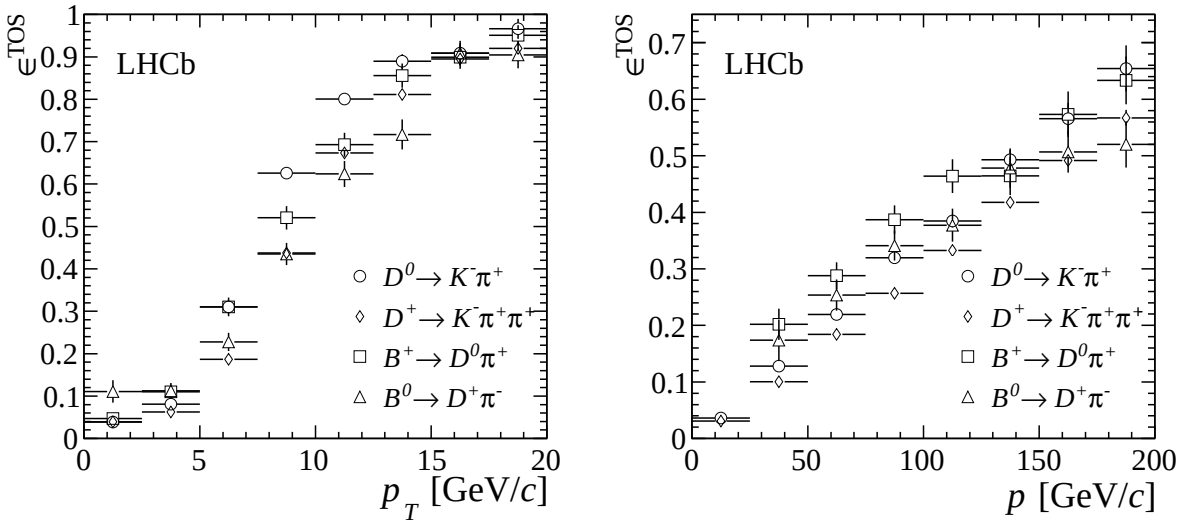

Figure 4. The efficiency $\varepsilon^{\rm TOS}$ of `L0Hadron` is shown for $B^0 \to D^+\pi^-$, $B^+ \to D^0\pi^-$, $D^0 \to K^-\pi^+$ and $D^+ \to K^-\pi^+\pi^+$ as a function of p_T and p of the signal.

`L0Electron` selects decays with electrons in the final state. It also triggers on radiative decays, with the photon being either converted, or with photon clusters with SPD hits in front due

to overlapping charged particles. For this reason the L0Photon and L0Electron thresholds are forced to be equal in the bandwidth division described above. The performance of L0Electron is shown in figure 5 for $B^0 \to J/\psi(e^+e^-)K^{*0}$ as a function of $p_T(J/\psi)$. Contrary to L0Muon, L0Electron is not fully efficient for J/ψ with large p_T. L0Electron only contains clusters with the same cell size, which limits its efficiency. In addition, off-line electron are corrected for bremsstrahlung, which is not the case in L0.

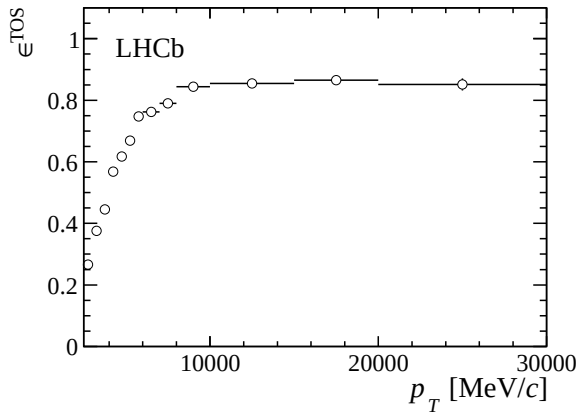

Figure 5. The efficiency ε^{TOS} of L0Electron is shown for $B^0 \to J/\psi(e^+e^-)K^{*0}$ as a function of p_T (J/ψ).

The performance of triggering on a photon is measured using $B^0 \to K^{*0}\gamma$ decays. The statistics in this channel for TIS events is insufficient to study the efficiency as a function of p_T or p, but allows to extract an efficiency of 50 ± 4 % for selecting the events with L0Photon alone. Selecting events with either L0Photon or L0Electron gives an efficiency of 88 ± 5 % .

Events selected by L0ElectronHi or L0PhotonHi are also selected by L0Electron or L0Photon, and hence do not contribute to the total L0 rate. These trigger lines are used in HLT1 to enhance the efficiency to select radiative B-decays, which is discussed in the next section.

5.3 HLT1 Performance

HLT1 muon trigger lines are only executed for events which have been triggered by L0Muon or L0DiMuon, and the lines require their tracks to be validated as a muon candidate as described in section 3.1. Table 8 gives the names of the HLT1 muon trigger lines and their cuts. Hlt1TrackMuon selects B, D or τ decays with at least one muon in its final state by requiring at least one muon with a significant IP. To trigger on muons originating from heavy particles with a negligible lifetime, like $W^{\pm}$, Z^0 or other heavy particles with a prompt muon in their final state, Hlt1SingleMuonHighPt has been implemented. It does not have any requirements on IP, but requires a hard p_T cut to reduce the rate. Hlt1DiMuonHighMass is complementary to Hlt1TrackMuon in that it allows selecting b-hadron decays without imposing lifetime related cuts, and thus allows to study the lifetime acceptance bias of the larger efficiency Hlt1TrackMuon line with high precision. Finally Hlt1DiMuonLowMass allows triggering on final states with two muons with a relative small invariant mass. To reduce the rate it imposes that the two muons are not prompt. The performance of the Hlt1 muon lines is presented using

Table 8. HLT1 muon trigger lines and their cuts. The rate is measured on events accepted by `L0Muon` or `L0DiMuon`.

Hlt1line	`TrackMuon`	`SingleMuon HighPT`	`DiMuon HighMass`	`DiMuon LowMass`
Track IP [mm]	> 0.1	-	-	-
Track IP χ^2	> 16	-	-	> 3
Track p_T [GeV/c]	> 1	> 4.8	> 0.5	> 0.5
Track p [GeV/c]	> 8	> 8	> 6	> 6
Track χ^2/ndf	< 2	< 4	< 4	< 4
DOCA [mm]	-	-	< 0.2	< 0.2
χ^2_{vertex}	-	-	< 25	< 25
Mass [GeV/c^2]	-	-	> 2.7	> 1
Rate [kHz]	5	0.7	1.2	1.3

$B^+ \to J/\psi K^+$. Figure 6 shows the performance of `Hlt1TrackMuon`, `Hlt1DiMuonHighMass` and `Hlt1DiMuonLowMass` as a function of the p_T, p and lifetime of the B^+. `Hlt1TrackMuon` gives the best performance overall, except at low lifetimes, where `Hlt1DiMuonHighMass` recovers events. `Hlt1DiMuonLowMass` looses ~ 10 % in efficiency compared to `Hlt1TrackMuon` for $B^+ \to J/\psi K^+$ due to the requirement to have at least two muon candidates, however it has much relaxed cuts on IP and p_T to allow the selection of candidates with the muon pair mass down to 1 GeV/c^2, which is useful for $B \to \mu\mu s$ decays.

The performance of `Hlt1SingleMuonHighPt` is not properly assessed using $B^+ \to J/\psi K^+$ decays because it is targeting decays of particles with a mass larger than b-hadrons. Instead $Z^0 \to \mu^+\mu^-$ events are used to measure the efficiency, by requiring one of the two muons to be TIS. This yields an efficiency of 77.1 ± 0.2 % for the `Hlt1SingleMuonHighPt` line per muon.

Apart from the above mentioned muon lines, HLT1 also contains a line which is executed for all L0 triggers, `Hlt1TrackAllL0`, which is designed to select hadron decays with a finite lifetime. A line called `Hlt1TrackPhoton` is only executed for events which have been triggered by `L0ElectronHi` or `L0PhotonHi`. This line is designed to enhance the trigger efficiency for radiative b-hadron decays with a high p_T photon. The corresponding selection cuts are given in table 9. Both trigger lines require at least one track with sufficient IP and p_T. `Hlt1TrackPhoton` aims at lower p_T tracks, and correspondingly also has relaxed track quality requirements compared to `Hlt1TrackAllL0`. figure7 shows the performance of `Hlt1TrackAllL0` as a function of p, p_T and τ for a few selected channels with hadronic decays. `Hlt1TrackAllL0` provides a very efficient trigger for all heavy flavour decays with a finite flight distance from their PV. At low p_T the requirement of at least one decay particle with $p_T > 1.7$ GeV/c results in selecting b-hadrons with a larger efficiency than c-hadrons, and low multiplicity decays with a larger efficiency than larger multiplicity decays. At large p_T this condition favours the decays with larger multiplicities.

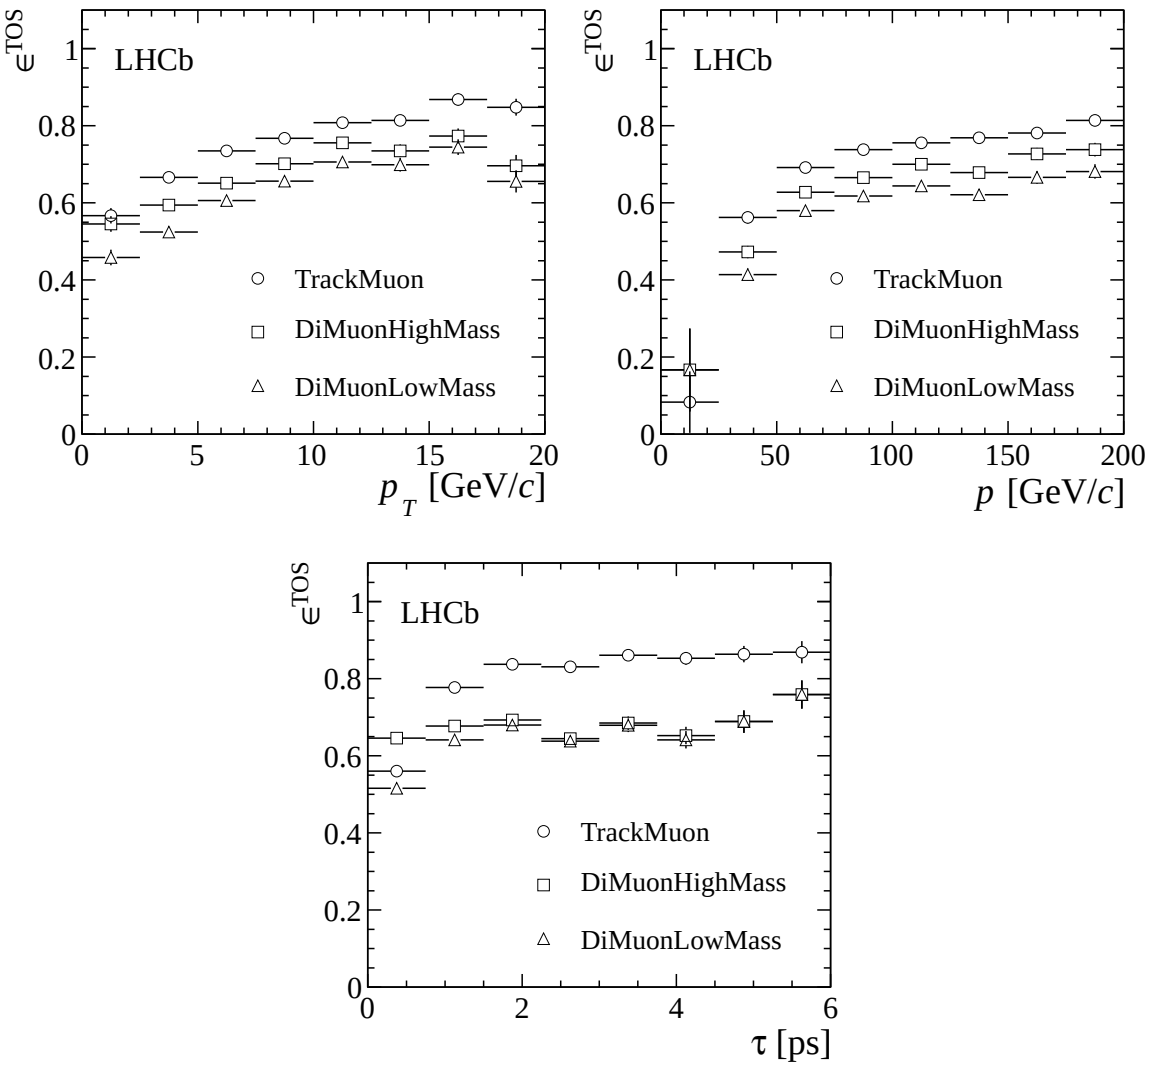

Figure 6. Efficiency $\varepsilon^{\rm TOS}$ of `Hlt1TrackMuon`, `Hlt1DiMuonHighMass` and `Hlt1DiMuonLowMass` for $B^+ \to J/\psi K^+$ as a function of the p_T, p and lifetime of the B^+.

While in L0 special trigger lines are used to trigger on b-hadron decays with electrons in the final state, in HLT1 these decays are covered by `Hlt1TrackAllL0`.

There is insufficient data in radiative B-decays to extract the performance of `Hlt1TrackPhoton` as a function of p_T in a data driven way. `Hlt1TrackPhoton` uses the same tracks as `Hlt1TrackAllL0`, but with a relaxed set of requirements as shown in table 9. The yield increase in $B^0 \to K^{*0}\gamma$ events obtained by including `Hlt1TrackPhoton` in addition to `Hlt1TrackAllL0` is measured to be 12 ± 2 %.

`Hlt1TrackAllL0` is complementary to trigger lines which require at least one of the decay muons to be identified. This is shown in Fig 8, which shows the performance of `Hlt1TrackMuon` for $B^+ \to J/\psi K^+$ decays, and the efficiency for the same channel when `Hlt1TrackAllL0` or `Hlt1TrackMuon` selected the event. The combined efficiency increases compared to the lines

Table 9. The cuts applied in `Hlt1TrackAllL0` and `Hlt1TrackPhoton` trigger lines. The rate is measured on events accepted by L0.

Hlt1line	`Hlt1TrackAllL0`	`Hlt1TrackPhoton`
Track IP [mm]	> 0.1	>0.1
Number VELO hits/track	> 9	> 6
Number missed VELO hits/track	< 3	< 3
Number OT+IT$\times 2$ hits/track	> 16	> 15
Track IPχ^2	> 16	>16
Track p_T [GeV/c]	> 1.7	> 1.2
Track p [GeV/c]	> 10	> 6
Track χ^2/ndf	< 2.5	< 2.5
Rate [kHz]	33	4.2

individually, even for decays with muons with relatively large p_T.

5.4 HLT2 Performance

Like in HLT1, HLT2 has lines which select events with one or two identified muons in the final state. In HLT2 the muon identification is identical to the off-line algorithm. The cuts corresponding to lines which are purely based on a single identified muon are given in table 10. `Hlt2SingleMuon` selects semileptonic b and c-hadron decays while avoiding a bias on the

Table 10. HLT2 trigger lines based on one identified muon.

`Hlt2Single`	`Muon`	`MuonHighPT`
`Hlt1TrackMuon`	TOS	-
Track IP [mm]	> 0.5	-
Track IPχ^2	> 200	-
Track p_T [GeV/c]	> 1.3	>10
Track χ^2/ndf	< 2	-
Pre-scale	0.5	1.
Rate [Hz]	480	45

hadronic part of the decay. To minimise the bias the p_T cut has been set low in combination with scaling the rate down by a factor two, rather than tightening the cut to reduce the rate. This line also provides a large yield for detached $J/\psi \to \mu\mu$ events which are selected by one of the two muons, while the other muon is used for calibration of tracking and muon identification effi-

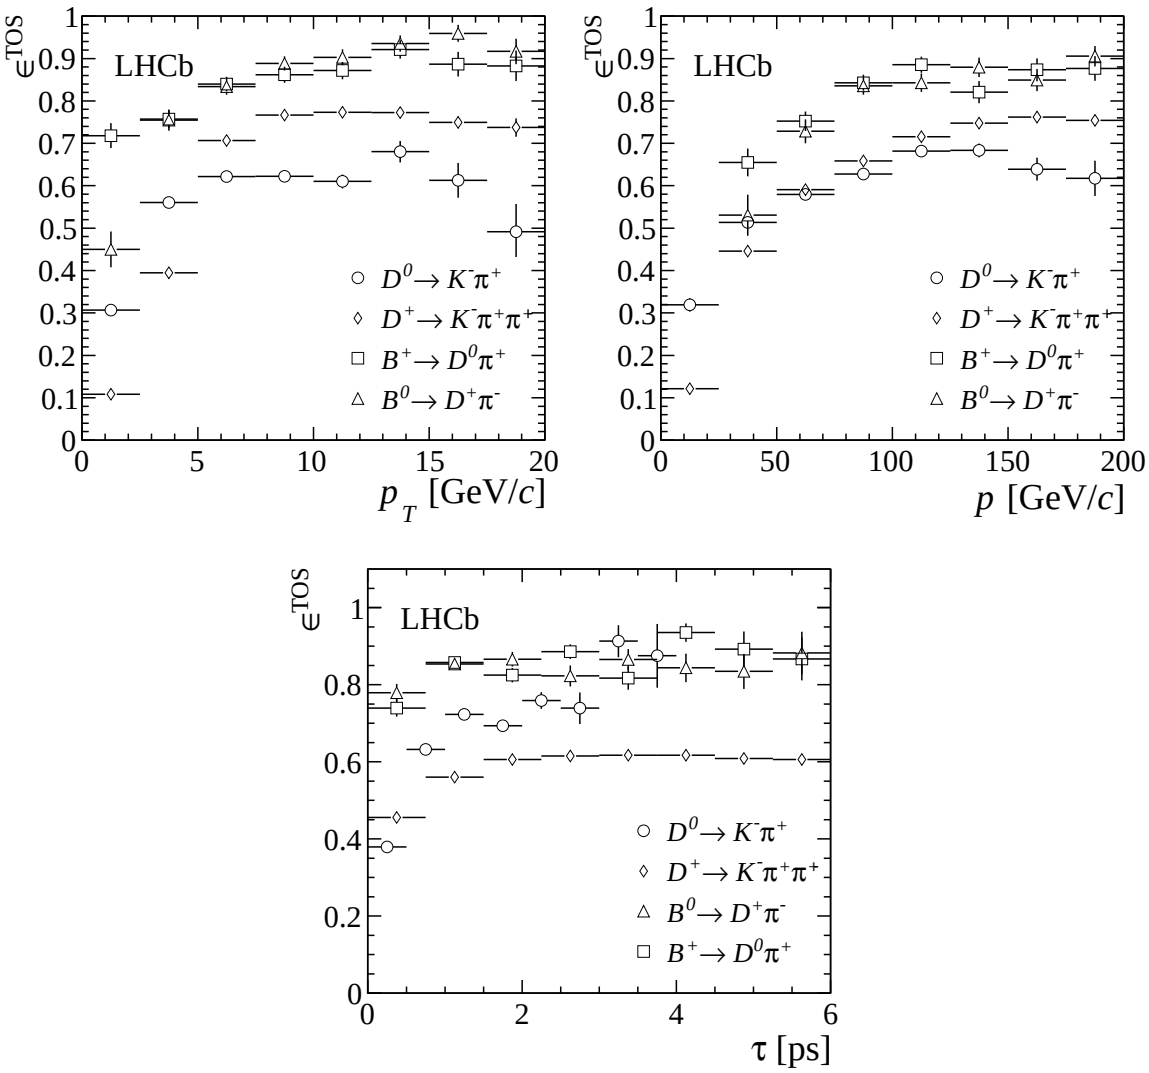

Figure 7. Efficiency ε^{TOS} of `Hlt1TrackAllL0` is shown for $D^0 \to K^-\pi^+$, $B^+ \to D^0\pi^-$ and $B^0 \to D^+\pi^-$ as a function of p_T, p and τ of the B-meson and prompt D-meson respectively.

ciencies. `Hlt2SingleMuonHighPT` provides the trigger for heavy particles decaying promptly to one or more muons, like $W^{\pm}$ or Z^0.

The Hlt2 triggers based on two identified muons distinguish lines with prompt decay selections, where the mass is the dominant discriminant, and lines which require a detached $\mu\mu$ vertex. The names and corresponding cuts of the prompt decay selections are given in Tables 11. `Hlt2DiMuonJPsi(Psi2S)` and `Hlt2DiMuonJPsi(Psi2S)HighPT` all select a mass region around J/ψ ($\psi(2S)$). `Hlt2DiMuonJPsi(Psi2S)` avoid explicit p_T requirements but as a consequence need to be pre-scaled to reduce their rate. `Hlt2DiMuonJPsi(Psi2S)HighPT` reduce the prompt J/ψ ($\psi(2S)$) rate by applying a p_T cut on the J/ψ ($\psi(2S)$) candidate. `Hlt2DiMuonB` has its mass cut set high enough to have an acceptable rate. The names and corresponding cuts of the detached decay selections are given in Tables 12. `Hlt2DiMuonDetached` is the main trig-

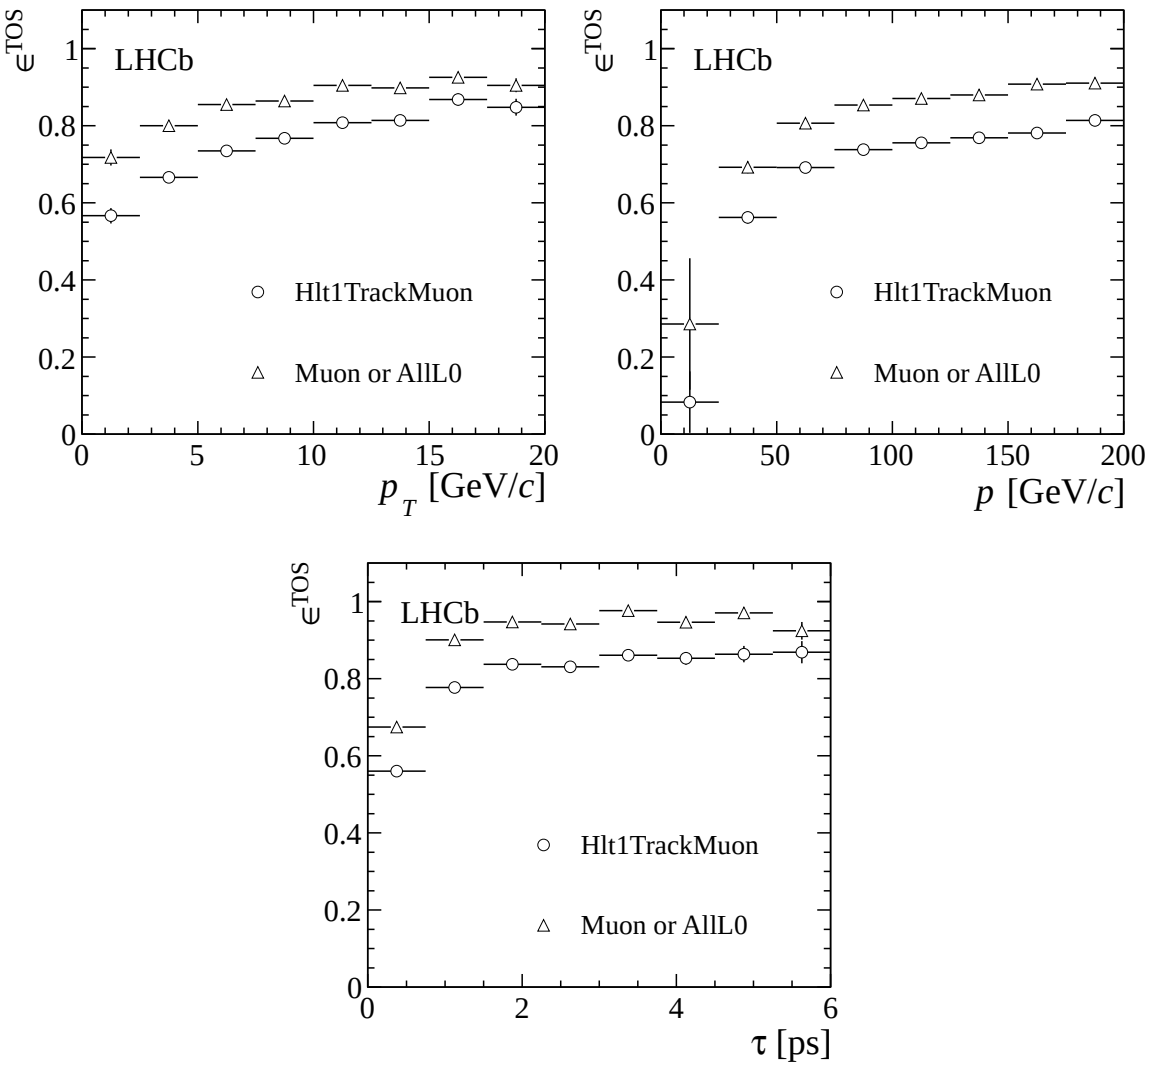

Figure 8. Efficiency ε^{TOS} of `Hlt1TrackMuon` is shown for $B^+ \to J/\psi K^+$ ($J/\psi \to \mu^+\mu^-$) as a function of p, p_T and of τ of the B-meson. Also shown is ε^{TOS} if at least one of the lines `Hlt1TrackAllL0` or `Hlt1TrackMuon` selected the events. The efficiency is measured relative to events which are TOS in `L0Muon` or `L0DiMuon`.

ger for low mass muon pairs. `Hlt2DiMuonDetachedHeavy` is an analogue line for J/ψ and higher mass muon pairs, with relaxed lifetime selection criteria. `Hlt2DiMuonDetachedJPsi` enhances the efficiency for J/ψ by reducing the flight distance requirement for J/ψ candidates even further. Figure 9 compares the performance of `Hlt2DiMuonJPsiHighPT` and `Hlt2DiMuonDetachedJPsi` as a function of the B p_T, p and τ in $B^+ \to J/\psi(\mu\mu)K^+$ decays, showing a very high efficiency for reconstructing and identifying muons in the trigger. `Hlt2DiMuonJPsi` avoids by design a bias in the proper lifetime, while the detached trigger lines allow to select decays with a significant flight distance with high efficiency at an affordable rate.

There are nine BBDT topological lines, `Hlt2Topo`*n*`Body`, `Hlt2TopoMu`*n*`Body` and `Hlt2TopoE`*n*`Body`,

Table 11. HLT2 trigger lines based on two identified muons without finite lifetime requirement.

Hlt2DiMuon	JPsi	Psi2S	B	JPsiHighPT	Psi2SHighPT
Track χ^2/ndf	< 5	< 5	< 5	< 5	< 5
Mass [GeV/c^2]	$M_{J/\psi} \pm 0.12$	$M_{\psi(2S)} \pm 0.12$	> 4.7	$M_{J/\psi} \pm 0.12$	$M_{\psi(2S)} \pm 0.12$
χ^2_{vertex}	< 25	< 25	< 10	< 25	< 25
$p_T{}^{\mu\mu}$ [GeV/c]	-	-	-	> 2	> 3.5
Pre-scale	0.2	0.1	1.	1.	1.
Rate [Hz]	50	5	80	115	15

Table 12. HLT2 trigger lines based on two identified muons with finite lifetime requirement.

Hlt2DiMuon	Detached	DetachedHeavy	DetachedJPsi
Track χ^2/ndf	< 5	< 5	< 5
Track IPχ^2	> 9	-	-
Mass [GeV/c^2]	> 1	> 2.95	$M_{J/\psi} \pm 0.12$
FDχ^2	>49	>25	>9
χ^2_{vertex}	< 25	< 25	< 25
$p_T{}^{\mu\mu}$ [GeV/c]	> 1.5	-	-
Rate [Hz]	70	75	35

where n=2,3,4 for the multiplicities considered, and TopoMu (TopoE) requiring at least one of the decay particles to have been identified as a muon (electron). Each line returns an output of the BBDT between 0 and 1. The Hlt2Topo$n$Body accept events with a rate of 930 Hz with a cut on the BBDT output at 0.4, 0.4 and 0.3 for the 2, 3 and 4 body lines respectively. While TopoMu and TopoE lines are based on the same BBDT, the extra requirement of either a muon or electron allows the cut on the BBDT output to be reduced to 0.1 for all six lines, which results in rates of 290 and 260 Hz for TopoMu and TopoE respectively.

The performance of the topological lines is given in figure 10 for hadronic B-decays and figure 11 for $B^+ \to J/\psi K^+$. Figure 11 also shows the complementarity of Hlt2Topo$n$Body and Hlt2TopoMunBody by showing the efficiency increase if either of these lines has selected the signal event. The inclusive performance of the topological lines is shown in figure 11 by giving the performance of Hlt2Topo2Body alone. This line requires only two of the three decay tracks of $B^+ \to J/\psi K^+$ to have been reconstructed and selected, and does not impose any particle identification. The combined performance of all topological lines mainly recovers efficiency at low p_T compared to the Hlt2Topo2Body line.

Table 5.4 gives the cuts applied in the two Hlt2 exclusive lines Hlt2CharmHadD02HH_D02KPi and Hlt2CharmHadD2HHH. The off-line selections of $D^0 \to K^-\pi^+$ are only slightly tighter than

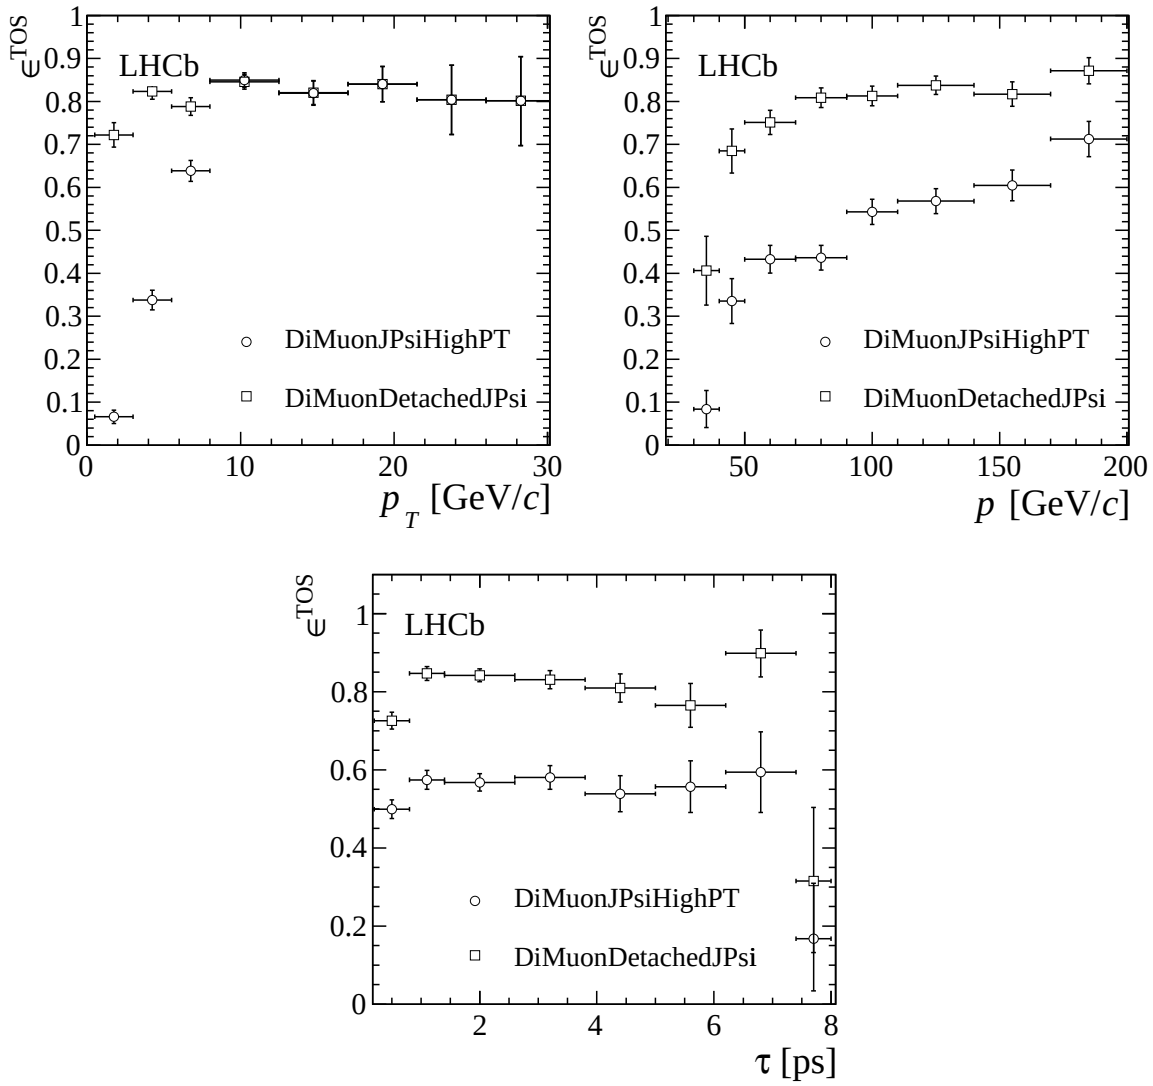

Figure 9. Efficiencies ε^{TOS} of `Hlt2DiMuonJPsiHighPT` and `Hlt2DiMuonDetachedJPsi` for $B^+ \to J/\psi K^+$ as a function of p_T, p and τ of the B^+.

the cuts applied in the trigger, resulting in an almost maximum efficiency of the `Hlt2CharmHadD02HH_D02KPi` line for this channel, as shown in figure 12. This figure also shows the performance of `Hlt2CharmHadD2HHH` for $D^+ \to K^-\pi^+\pi^+$. Here the trigger looses efficiency due to the necessity to first having to apply cuts to two of the three decay products as described in 3.2.3.

6. Summary

The LHCb trigger is designed to select charm and beauty hadrons for a large range of decay modes, and has the ability to measure the selection efficiency from the data. In 2011 the trigger has been tuned to cope with pp-interactions at $\sqrt{s} = 7$ TeV, with 1296 colliding bunches in LHCb with an

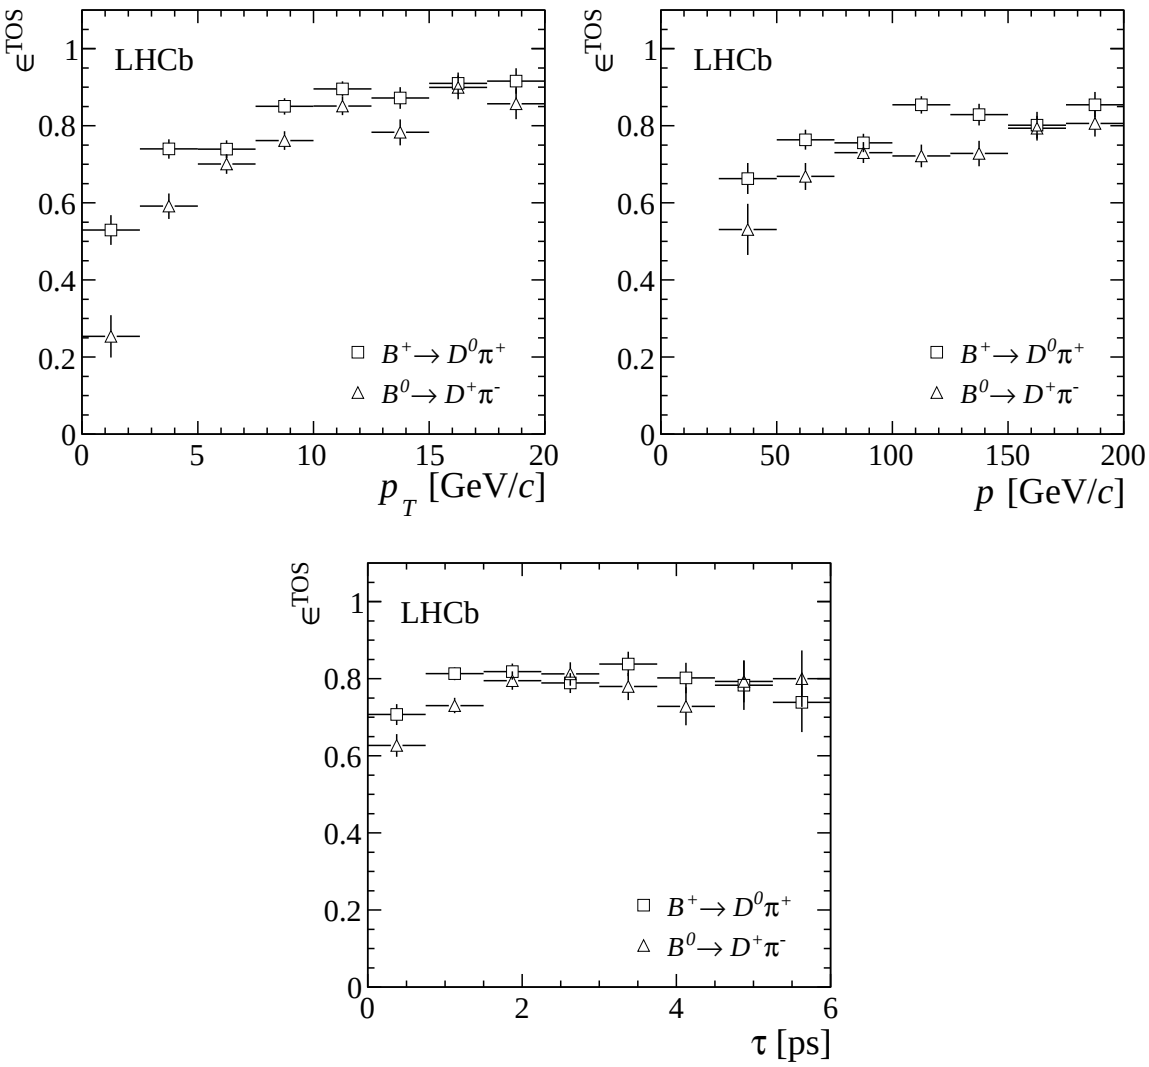

Figure 10. Efficiency ε^{TOS} if at least one of the lines `Hlt2Topo`n`Body`, with $n = 2,3$, selected the event for $B^+ \to D^0\pi^-$ and one of the lines with $n = 2,3,4$ for $B^0 \to D^+\pi^-$ as a function of p_T, p and τ of the B-meson. This efficiency is measured relative to events which are TOS in `Hlt1TrackAllL0`.

average number of visible pp-interactions of 1.4. This corresponds to a bunch crossing rate with at least one visible pp-interaction of ~ 11 MHz.

L0 reduced this rate to 870 kHz by applying p_T cuts on muons and E_T cuts on clusters in the calorimeters. HLT1 performs a partial reconstruction of tracks and performs lepton identification. It employs a combination of cuts on p_T, invariant mass and IP to reduce the rate to around 43 kHz. HLT2 reconstructs all tracks in the event with $p_T >$500 MeV/c. It selects candidates using as signatures leptonic decays, finite lifetime and invariant mass. Its output rate is 3 kHz, which is subdivided in 50 % inclusive hadronic triggers, 25 % triggers on leptons and the remaining rate from exclusive triggers, mainly on charmed hadrons.

The efficiency for the major trigger lines have been presented for representative decay modes

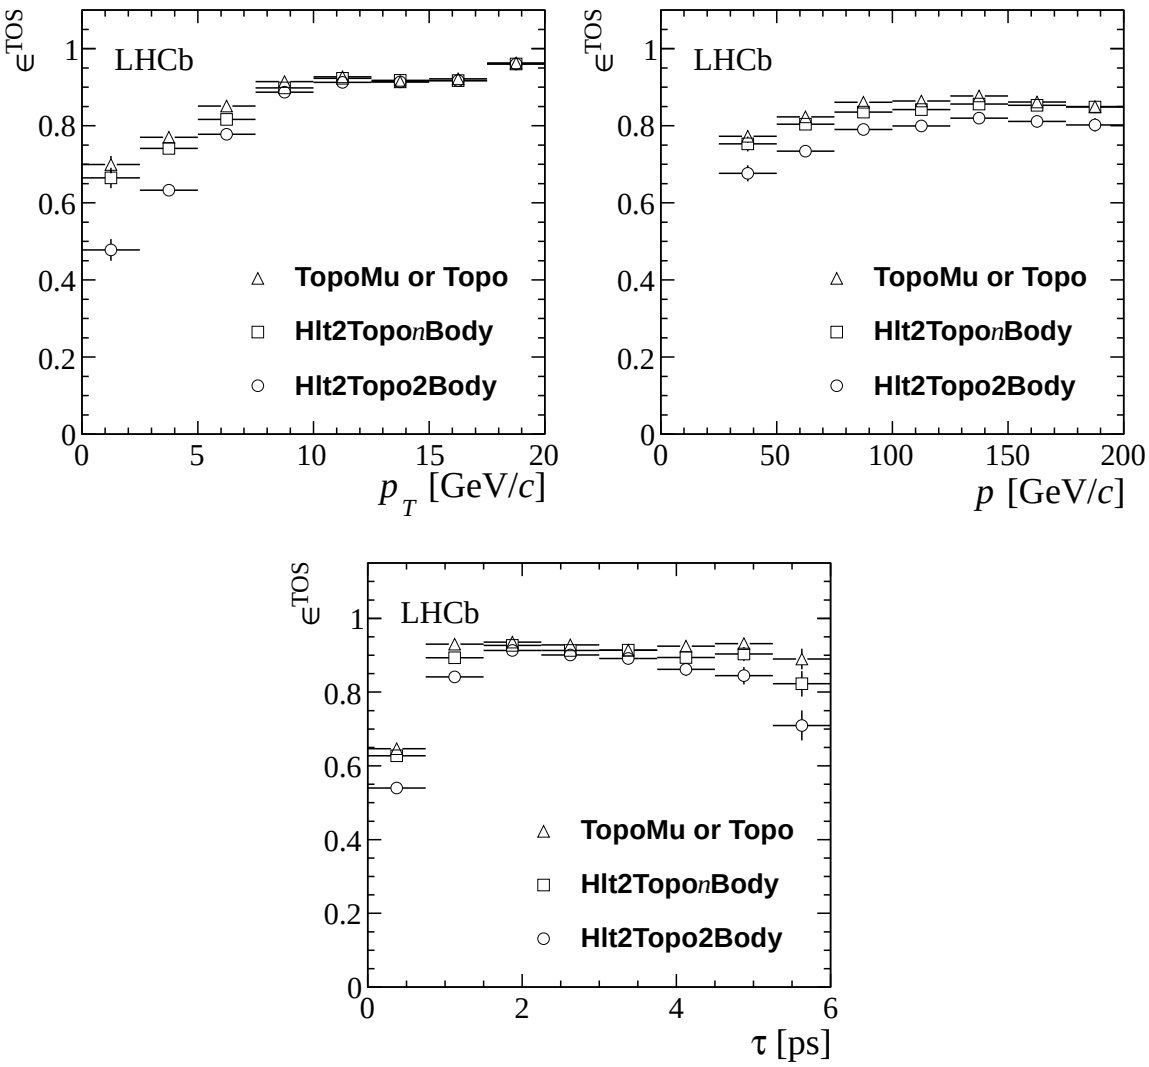

Figure 11. Efficiency ε^{TOS} if at least one of the lines `Hlt2Topo`n`Body` and `Hlt2TopoMu`n`Body`, with $n = 2, 3$, selected events for $B^+ \to J/\psi K^+$, as a function of B p_T, p and τ. Also shown is ε^{TOS} if at least one of the lines `Hlt2Topo`n`Body` or `Hlt2TopoMu`n`Body` selected the events. `Hlt2Topo2Body` shows the inclusive performance of the topological lines. Efficiency is measured relative to events which are TOS in either `Hlt1TrackAllL0` or `Hlt1TrackMuon`.

as a function of p_T, p and lifetime of c and b-hadrons.

Acknowledgments

We express our gratitude to our colleagues in the CERN accelerator departments for the excellent performance of the LHC. We thank the technical and administrative staff at the LHCb institutes. We acknowledge support from CERN and from the national agencies: CAPES, CNPq, FAPERJ and FINEP (Brazil); NSFC (China); CNRS/IN2P3 and Region Auvergne (France); BMBF, DFG, HGF and MPG (Germany); SFI (Ireland); INFN (Italy); FOM and NWO (The Netherlands); SCSR

Table 13. Hlt2 selection cuts applied for the exclusive lines `Hlt2CharmHadD02HH_D02KPi` and `Hlt2CharmHadD2HHH`. The 2-track cuts refer to a candidate constructed of two tracks, and $m_{\rm corr}$ is defined in equation 2. The angle α is the angle between the momentum of the B and the vector connecting the PV with the D vertex.

Variable	`Hlt2CharmHadD02HH_D02KPi`	`Hlt2CharmHadD2HHH`
$\chi^2_{\rm track}$/ndf	< 3	<3
Track p_T [MeV/c]	> 800	>250
Track p [MeV/c]	> 5000	>2000
track Σp_T [MeV/c]	-	> 2500
$\geqq 1$ track p_T [MeV/c]	> 1500	-
2-track mass [MeV/c^2]	-	< 2100
2-track $m_{\rm corr}$ [MeV/c^2]	-	<3500
2-track IPχ^2	-	>40
$\geqq 2$ tracks p_T [MeV/c]	-	>500
$\geqq 2$ tracks p [MeV/c]	-	>5000
Tracks IPχ^2	> 9	> 5
$\geqq 2$ tracks IPχ^2	> 9	> 10
2-track DOCA [mm]	<0.1	<0.1
$\chi^2_{\rm vertex}$/ndf	< 10	< 20
FDχ^2	> 40	> 150
D IPχ^2	-	< 12
D cos(α)	> 0.99985	-
D p_T [MeV/c]	> 2000	> 1000
D mass interval [MeV/c^2]	1815-1915	1800-2040
Rate [Hz]	260	390

(Poland); ANCS/IFA (Romania); MinES, Rosatom, RFBR and NRC “Kurchatov Institute” (Russia); MinECo, XuntaGal and GENCAT (Spain); SNSF and SER (Switzerland); NAS Ukraine (Ukraine); STFC (United Kingdom); NSF (USA). We also acknowledge the support received from the ERC under FP7. The Tier1 computing centres are supported by IN2P3 (France), KIT and BMBF (Germany), INFN (Italy), NWO and SURF (The Netherlands), PIC (Spain), GridPP (United Kingdom). We are thankful for the computing resources put at our disposal by Yandex LLC (Russia), as well as to the communities behind the multiple open source software packages that we depend on. A special acknowledgement goes to all our LHCb collaborators who over the years have contributed to obtain the results presented in this paper.

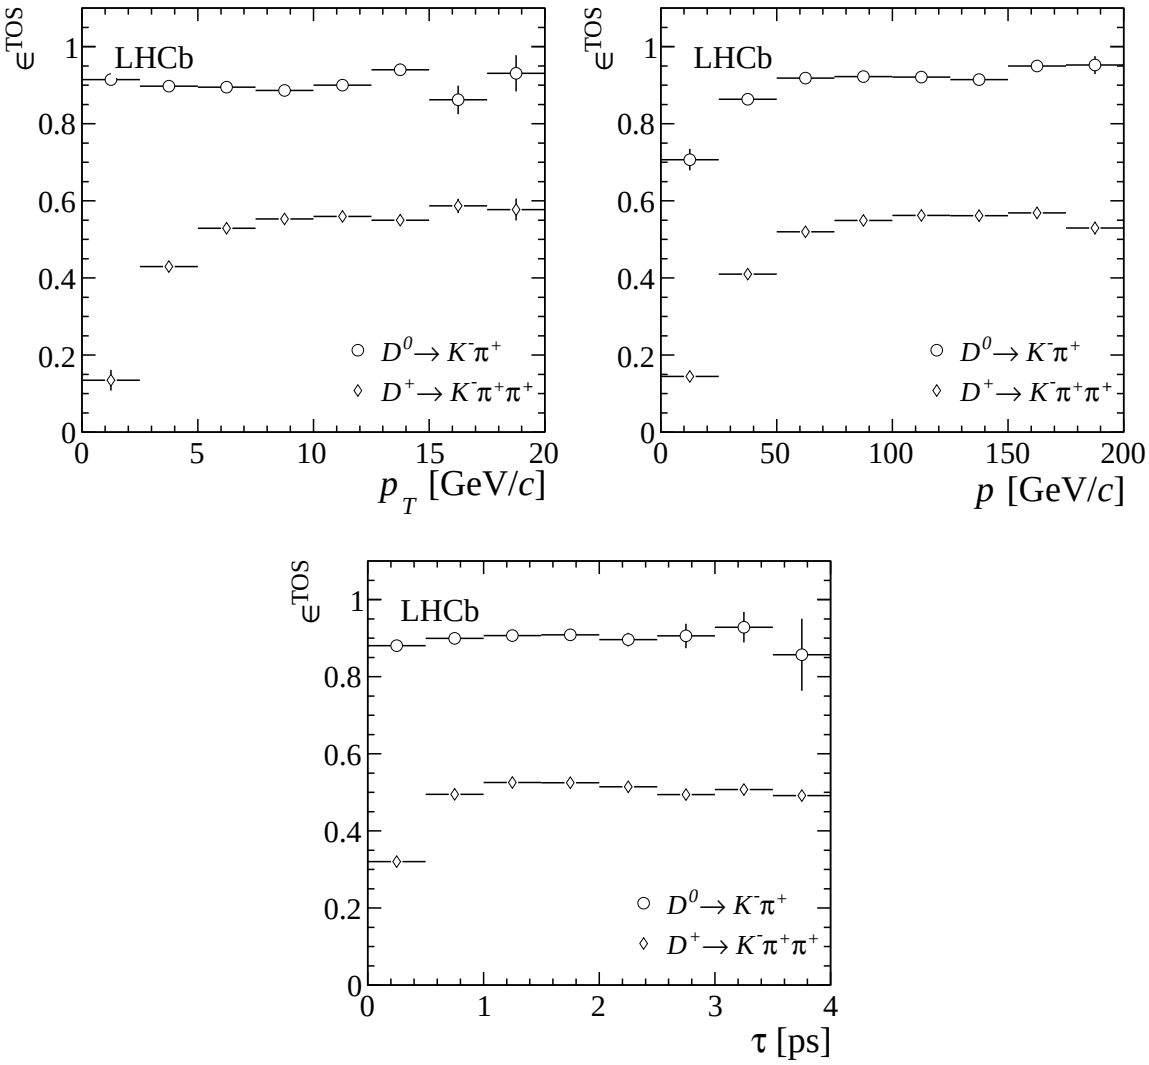

Figure 12. Efficiency ε^{TOS} of the lines `Hlt2CharmHadD2HHH` and `Hlt2CharmHadD02HH_D02KPi` for $D^+ \to K^-\pi^+\pi^+$ and $D^0 \to K^-\pi^+$ respectively as a function of p_T, p and τ. This efficiency is measured relative to events which are TOS in `Hlt1TrackAllL0`.

References

[1] The LHCb Collaboration, *The LHCb Detector at the LHC*, 2008 *JINST* **3** S08005

[2] G. Corti et al., *Software for the LHCb experiment*, IEEE Trans. Nucl. Sci. **53** (2006) 1323

[3] The LHCb Collaboration, *Roadmap for selected key measurements of LHCb*, `hep-ex/0912.4179v3`

[4] The LHCb Collaboration, *Absolute luminosity measurements with the LHCb detector at the LHC*, 2012 *JINST* **7** P01010

[5] R.E. Kalman, *A new approach to linear filtering and prediction problems*, Trans. ASME J. Bas. Eng. D82 (1960) 35.

R. Frühwirth, *Application of Kalman Filtering to track and vertex fitting*. Nucl. Instrum. Meth. A 262 (1987) 444.

[6] M. Gandelman and E. Polycarpo, *The Performance of the LHCb Muon Identification Procedure*, LHCb-2007-145. CERN-LHCb-2007-145.

[7] SLD Collaboration, *Measurement of Rb using a Vertex Mass Tag*, Phys. Rev. Lett., 80:660-665 (1998).

[8] L. Brieman, J. H. Friedman, R. Olshen, and C. J. Stone, *Classication and Regression Trees*, Wadsworth International Group, Belmont, California, 1984.
L. Breiman, *Bagging predictors*, Machine Learning 24 (1996) 123.
H. Yang, B.P. Roe, J. Zhu., *Studies of Boosted Decision Trees for MiniBooNE Particle Identification*, arXiv:physics/0508045v1

[9] M. Feindt and U. Kerzel, *The NeuroBayes neural network package*, Nucl. Instrum. Meth. A **559** (2006) 190.